

TANGLE STRUCTURE TREES II: TREES OF TANGLES AND TANGLE-TREE DUALITY

HANNO VON BERGEN AND REINHARD DIESTEL

ABSTRACT. Tangle structure trees, introduced in [3], offer a unified data structure that displays all the tangles of a graph or data set together with certificates for the non-existence of any other tangles, either locally or overall. In this paper we apply tangle structure trees to derive new versions of the two fundamental tangle theorems: the tree-of-tangles theorem, and the tangle-tree duality theorem.

We extend the tree-of-tangles theorem to $\mathcal{F}$ -tangles that need not be profiles. When $\mathcal{F}$ consists of stars of separations, as it does in classical tangle-tree duality theorems, we show how to convert tangle structure trees that certify the non-existence of $\mathcal{F}$ -tangles into tree-decompositions that certify this in the way known from graph tangles, as S -trees over $\mathcal{F}$.

1. INTRODUCTION

In this sequel to [3] we explore further applications of *tangle structure trees*. These are data structures, introduced in [3], which simultaneously display

- all the tangles of a given graph or abstract separation system; and
- for all its orientations that are not tangles, certificates of why they are not.

Previously, there were two separate types of theorem in tangle theory for these two purposes: *tree-of-tangle* theorems, which display all the tangles of a graph or data set in a tree-like way that ‘distinguishes’ different tangles by specifying a minimum-order separation that they orient differently; and *tangle-tree duality* theorems that display, also in a tree-like way, certificates for why the various orientations of the separation system are not tangles – typically in the form of small subsets of oriented separations that are *forbidden* in the given type of tangle. These two fundamental theorems, which exist in various forms and degrees of generality, are thought of as the two pillars of both graph and abstract tangle theory.

The unified way in which tangle structure trees serve both these purposes still allows us to extract, if desired, separate corollaries of the two types. In [3] we did this for the second type: we derived new tangle-tree duality-type theorems for tangles whose defining forbidden subsets were not ‘stars’ of separations, nested sets of typically three oriented separations pointing towards each other. All previously known tangle-tree duality theorems, in particular those for graphs, were for tangles defined by excluding such stars.

Our first main result in this paper is to extract from the tangle structure trees established in [3] a new theorem of the first type: a tree-of-tangles theorem for tangles whose set $\mathcal{F}$ of defining forbidden subsets need not form a ‘profile’. All known tree-of-tangles theorems so far are for profiles, in fact, for *robust* profiles.

These are still a sweeping generalization of graph tangles. But we shall prove here that we can generalize them further: we shall obtain a tree-of-tangles theorem for all robust $\mathcal{F}$ -tangles, no matter what $\mathcal{F}$ is.

Our second main result is about $\mathcal{F}$ -tangles whose set $\mathcal{F}$ of forbidden subsets consists of stars of separations, as in all the classical tangle-tree duality theorems. We prove a general conversion theorem, which allows us to extract from our tangle structure tree for $\mathcal{F}$ -tangles of a separation system $\vec{S}$, say, an S -tree over $\mathcal{F}$ if S admits no $\mathcal{F}$ -tangles. Such an ‘ S -tree over $\mathcal{F}$ ’ is the tree structure in the most general tangle-tree duality theorems known so far for abstract separation systems (which include graphs), proved in [11]. Our conversion theorem allows us to reprove the main result of [11] from the premise for the existence of tangle structure trees, which differs from the premise of the tangle-tree duality theorem in [11].

However we show that our premise for the existence of tangle structure trees follows from an only slightly stronger version of the premise in [11], the version used in practice when the duality theorem from [11] is applied in [10]. As a consequence, we shall be able in [2] to use our conversion theorem, coupled with the existence theorems for tangle structure trees in [3], to derive all the known applications of the tangle-tree duality theorem in [11] to concrete structures, such as graphs, matroids, or data sets [10].

The paper is organized as follows. In [Section 2](#) we give a summary of tangle basics, just what is technically needed to read this paper. More on graph tangles can be found in [7], more on abstract tangles and their applications in [4–6]. In [Section 3](#) we introduce tangle structure trees and summarize the main results from [3]; see there more motivational background for tangle structure trees.

[Section 4](#) contains our conversion theorem that extracts from a tangle structure tree of a separation system $\vec{S}$ that has no $\mathcal{F}$ -tangles an S -tree over $\mathcal{F}$. [Section 5](#) is an interlude on submodular order functions for separations, which is needed in the rest of the paper; we show how a non-injective submodular order function can be extended to an injective one without losing submodularity. The tree-of-tangles theorem for arbitrary robust $\mathcal{F}$ -tangles is proved in [Section 6](#). In [Section 7](#) we show how to recover the essence of the tangle-tree duality theorem of [11] from our conversion theorem and our existence theorem for tangle structure trees.

2. SEPARATION SYSTEMS AND THEIR TANGLES

In this section we introduce the terminology that abstract tangle theory requires, following [4].¹ Readers familiar with [3] may skip this section.

Tangles of graphs are ways of orienting their separations, each towards one of its two sides. Abstract tangles are designed to work in scenarios where there need not be anything to ‘separate’. In order to retain our intuition from graphs, however, we continue to refer to the things of which our abstract tangles pick one of two variants (which they will indeed do) as ‘separations’. These are defined by noting some key properties of graph separations and making them into axioms, as follows.

A *separation system* $(\vec{S}, \leq, *)$ is a set $\vec{S}$, whose elements we call *oriented separations*, that comes with a partial ordering $\leq$ on $\vec{S}$ and an order-reversing involution

¹There is one difference to [4]: for historical reasons, the partial ordering on $\vec{S}$ used there is the inverse of ours. So terms like ‘large’ and ‘small’, infima and suprema etc, are reversed.

$^*: \vec{S} \rightarrow \vec{S}$. Thus, for any two elements² $\vec{r}, \vec{s}$ of $\vec{S}$ with $\vec{r} \leq \vec{s}$ we have $\vec{r}^* \geq \vec{s}^*$. We write $\vec{s}^* =: \vec{s}$, and call $\vec{s}$ the *inverse* of $\vec{s}$. While we allow formally that $\vec{s} = \vec{s}$, in which case we call $\vec{s}$ and s *degenerate*, this does not happen often in practice.³

If a separation system $\vec{U}$ happens to be a lattice, that is, if there is a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$ for every two elements $\vec{r}, \vec{s} \in \vec{U}$, we call $\vec{U}$ a *universe* of separations. It is *distributive* if it is distributive as a lattice. A separation system $\vec{S} \subseteq \vec{U}$ is *submodular* if for every two elements of $\vec{S}$ either their infimum or their supremum in $\vec{U}$ also lies in $\vec{S}$.

Very rarely we may have separations $\vec{s} \leq \vec{s}$; then $\vec{s}$ is *small* and $\vec{s}$ is *large*.⁴ Separation systems without small elements are *regular*. We say that $\vec{s}$ is *trivial* (and $\vec{s}$ is *co-trivial*) in $\vec{S}$ if there exists a pair of inverse separations $\vec{r}, \vec{r} < \vec{s}$ in $\vec{S}$. Trivial separations are clearly large, so co-trivial ones are small, but the converse need not hold. See [4] for more on these technicalities if desired.

The set of *unoriented separations* in $(\vec{S}, \leq, ^*)$ is

$$S := \{\{\vec{s}, \vec{s}\} : \vec{s} \in \vec{S}\}.$$

We call the elements $\vec{s}, \vec{s}$ of s its *orientations*. An *orientation* of S is a set $\tau \subseteq \vec{S}$ that contains exactly one orientation of every $s \in S$. An orientation of a subset of S is a *partial orientation* of S . We write $\tau(s)$ for the unique orientation of s contained in τ , and say that s *distinguishes* two partial orientations τ, τ' of S if both are defined on s and $\tau(s) \neq \tau'(s)$.

If $\vec{r} \geq \vec{s}$ we say that $\vec{r}$ *points towards* s (and that $\vec{r}$ *points away from* s). We say that $\vec{r}$ *points towards* an oriented separation $\vec{s}$ whenever it points towards s , i.e., if $\vec{r} \geq \vec{s}$ or $\vec{r} \geq \vec{s}$, and similarly for ‘points away from’. A *star* is a set σ of non-degenerate oriented separations that point towards each other. As is easy to check, this happens if and only if $\vec{r} \geq \vec{s}$ (and hence $\vec{s} \geq \vec{r}$) for all distinct $\vec{r}, \vec{s} \in \sigma$.




FIGURE 1. Nested separations $r = \{A, B\}$ and $s = \{C, D\}$ of a graph. Their orientations $\vec{r} = (A, B)$ and $\vec{s} = (D, C)$ point towards each other, since $\vec{r} \geq \vec{s}$ (as $B \supseteq D$) and $\vec{s} \geq \vec{r}$ (as $C \supseteq A$).

Two separations $r, s \in S$ are *nested* if they have orientations that are comparable under $\leq$; otherwise they *cross*. Oriented separations are *nested* if their underlying unoriented separations are nested. A subset of $\vec{S}$ is *nested* if its elements are pairwise nested.

²We often denote the elements of $\vec{S}$ by letters with an arrow, in either direction, precisely in order to have a simple way to refer to their dual elements: by reversing the arrow. But the arrow directions have no meaning: an arbitrary element of $\vec{S}$ could be denoted equally as $\vec{s}$ or as $\vec{s}$.

³The only degenerate separation of a graph $G = (V, E)$, for example, is $\{V, V\}$.

⁴The small separations of a graph G are those of the form (V, A) with $A \subseteq V = V(G)$.

A subset of $\vec{S}$ is *consistent* if no pair of its elements $\vec{r}, \vec{s}$ with $r \neq s$ point away from each other. Stars are examples of consistent nested sets of oriented separations.

We shall often be interested in consistent orientations of S . For each of its elements $\vec{r}$, a consistent orientation of S will also contain every $\vec{s} > \vec{r}$ other than, possibly, $\vec{s} = \vec{r}$. Consistent partial orientations of S are easily seen to extend to consistent orientations of S , unless they contain a separation that is co-trivial in $\vec{S}$; see [4].

If $\sigma \subseteq \vec{S}$ is consistent, we say that $\vec{s} \in \vec{S}$ is *required* by σ if $\vec{s} \notin \sigma$ and $\sigma \cup \{\vec{s}\}$ is inconsistent. We shall see in [Lemma 2.1](#) that, pathological cases aside, $\sigma \cup \{\vec{s}\}$ will then be consistent. The *closure* of σ is

$$[\sigma] := \sigma \cup \{\vec{s} \in \vec{S} : \vec{s} \text{ is required by } \sigma\}.$$

Note that σ requires $\vec{s} \notin \sigma$ if and only if there exists an $\vec{r} \in \sigma$ such that $r \neq s$ and $\vec{r} < \vec{s}$. Thus

$$[\sigma] = \sigma \cup \{\vec{s} \in \vec{S} : \exists \vec{r} \in \sigma \text{ such that } r \neq s \text{ and } \vec{r} < \vec{s}\},$$

which motivates the term of (upward) ‘closure’. If $[\sigma]$ is consistent, then $[[\sigma]]$ is (defined and) easily shown to be equal to $[\sigma]$.

If σ contains no small separations, the expression above simplifies to

$$[\sigma] = \{\vec{s} \in \vec{S} : \exists \vec{r} \in \sigma \text{ such that } \vec{r} \leq \vec{s}\}.$$

Indeed, any $\vec{s}$ in this latter set either lies in σ or there is some $\vec{r} \in \sigma$ such that $\vec{r} < \vec{s}$; in that case $r \neq s$, since otherwise $\vec{s} = \vec{r} < \vec{s}$, making $\vec{r} \in \sigma$ small.

Lemma 2.1 ([3]). *Let $\sigma \subseteq \tau \subseteq \vec{S}$ be consistent sets. Then $[\sigma] \subseteq [\tau]$. If τ is an orientation of all of S , then $[\tau] = \tau$. If τ has no elements that are co-trivial in $\vec{S}$, then $[\tau]$ is consistent and $[\tau] \setminus \tau$ contains at most one orientation of any $s \in S$.*

Consistent orientations of S can contain small separations. But examples are rare and can be counter-intuitive, so we often exclude them. Co-trivial separations cannot occur in consistent orientations of S . Indeed if $\vec{s}$ is co-trivial, witnessed by $r \in S$, then every orientation of S will have to orient r too. But it cannot do so consistently with $\vec{s}$, since both $\vec{r}$ and $\bar{r}$ are inconsistent with $\vec{s}$. Similarly, $[\{\vec{s}\}]$ is then inconsistent, since it contains $\vec{r}$ and $\bar{r}$, which are both inconsistent with $\vec{s}$.

An *order function* on S is any map $S \rightarrow \mathbb{R}$. Unless otherwise mentioned, we denote such order functions as $s \mapsto |s|$. We extend them to $\vec{S}$ by letting $|\vec{s}| := |\bar{s}| := |s|$. A separation system $(\vec{S}, \leq, *)$ given with an order function on S is an *ordered separation system*. For every $k \in \mathbb{R}$ we let

$$\vec{S}_k := \{\vec{s} \in \vec{S} : |s| < k\};$$

this is again an ordered separation system.

We are sometimes interested in orientations of S that do not have certain subsets. We typically collect those together in some set $\mathcal{F}$, whose elements we call *forbidden subsets*. Formally, if $\mathcal{F}$ is any set, we say that $\tau \subseteq \vec{S}$ *avoids* $\mathcal{F}$ if τ has no subset in $\mathcal{F}$, i.e., if no subset of τ is an element of $\mathcal{F}$.

Definition 2.2. An $\mathcal{F}$ -*tangle* of S is an $\mathcal{F}$ -avoiding consistent orientation of S . The $\mathcal{F}$ -tangles of the subsets S_k of S are the $\mathcal{F}$ -tangles *in* $\vec{S}$.

When our choice of $\mathcal{F}$ is clear from the context, or arbitrary, we shall also refer to $\mathcal{F}$ -tangles simply as *tangles*.

The tangles of the sets $S_k \subseteq S$ are the k -tangles of S , or the *tangles of order k* in $\vec{S}$. When we speak of *maximal* tangles in $\vec{S}$, we refer to their partial ordering as subsets of $\vec{S}$. A maximal $\mathcal{F}$ -tangle in $\vec{S}$, thus, is an $\mathcal{F}$ -tangle τ of some S_k for which there is no $\mathcal{F}$ -tangle $\tau' \neq \tau$ of any S_ℓ such that $\tau = \tau' \cap \vec{S}_k$.

3. TANGLE STRUCTURE TREES

We adopt the graph-theoretic terminology of [7]. Given a tree T with a root r , we write $\leq_r$ for the associated partial ordering on the nodes of T . Maximal elements, including the root if $|T| = 1$, are called *leaves*. Any direct successors in $\leq_r$ of a node of T are its *children*. We write E_v for the set of edges of T from a node v to its children.

Let $(\vec{S}, \leq_r^*)$ be a separation system, and let $\mathcal{F}$ be any set. A *separation tree* (T, r, β) on $\vec{S}$ consists of a rooted tree (T, r) and an edge labelling $\beta: E(T) \rightarrow \vec{S}$ such that for every non-leaf $v \in V(T)$ there exists a separation $s_v \in S$ such that β restricts to a bijection $E_v \rightarrow \{\vec{s}_v, \vec{s}_v^*\}$ and $s_u \neq s_v$ whenever $u <_r v$.

Thus, every non-leaf node v of a separation tree has either one or two children. If there is no need to refer to r or β explicitly, we usually abbreviate (T, r, β) to T . For every node v , we write β_v for the set $\beta(E(rTv))$ of edge labels on the path rTv . A separation tree is *consistent* if $\beta_v \subseteq \vec{S}$ is consistent for every node $v \in T$.

A leaf ℓ of a consistent separation tree on $\vec{S}$ is an *$\mathcal{F}$ -tangle leaf* if the closure $[\beta_\ell]$ of β_ℓ is an $\mathcal{F}$ -tangle of S . A non-leaf node v is an *$\mathcal{F}$ -tangle node* if there is an $\mathcal{F}$ -tangle leaf $\ell \geq_r w$ for every (one or two) successor w of v in $<_r$. Tangle nodes with two children, thus, are precisely the $<_r$ -infima of pairs of distinct tangle leaves.

A leaf ℓ is *forbidden* (by $\mathcal{F}$) if β_ℓ contains an element of $\mathcal{F}$ as a subset. Note that tangle leaves are never forbidden. An *$\mathcal{F}$ -tangle structure tree of $\vec{S}$* is a consistent separation tree on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf node v the set β_v has no subset in $\mathcal{F}$.

Theorem 3.1 ([3, Theorem 3.7]). *Let $\vec{S}$ be a separation system, let $\mathcal{F}$ be any set, and let T be an $\mathcal{F}$ -tangle structure tree of $\vec{S}$. Then for every $\mathcal{F}$ -tangle τ of S there is a unique leaf $\ell =: \ell(\tau)$ of T with $[\beta_\ell] = \tau$. In particular, if all the leaves of T are forbidden, then S has no $\mathcal{F}$ -tangle.*

$\mathcal{F}$ -tangle structure trees will be our main tool, so let us see when they exist.

Essentially, two conditions on $\mathcal{F}$ are needed to ensure that $\vec{S}$ has an $\mathcal{F}$ -tangle structure tree. The first is that $\{\vec{s}\} \in \mathcal{F}$ for every $\vec{s} \in \vec{S}$ that is trivial in $\vec{S}$. This holds for all relevant $\mathcal{F}$ (see [3] for why), and if it does we call $\mathcal{F}$ *standard for $\vec{S}$* .

The second condition is the one that bites. If T is any separation tree on $\vec{S}$, then every orientation τ of S contains a set β_ℓ for a unique leaf ℓ of T [3, Lemma 3.2]. If τ is consistent and $[\beta_\ell]$ orients all of S , then $[\beta_\ell] = \tau$ by Lemma 2.1. Hence if τ is not an $\mathcal{F}$ -tangle, then $[\beta_\ell]$ has a subset σ in $\mathcal{F}$ that witnesses this.

If T is even an $\mathcal{F}$ -tangle structure tree, then ℓ is a forbidden leaf, and so we can find such a subset $\sigma \in \mathcal{F}$ of τ not just in $[\beta_\ell]$ but even in β_ℓ : among the edge labels of T . In order for this to be possible, we therefore need to make some ‘richness’ assumption about $\mathcal{F}$: an assumption which, in our example, ensures that $\mathcal{F}$ has enough elements to contain a subset also of β_ℓ as soon as it contains a subset of $[\beta_\ell]$.

The main contribution of [3] was to identify such a richness condition on $\mathcal{F}$ that is, essentially, both necessary and sufficient for $\mathcal{F}$ -tangle structure trees to exist.

To motivate it, let us first state a stronger condition that is less technical, and which clearly implies that $\mathcal{F}$ contains a subset of β_ℓ whenever it contains a subset of $\lfloor \beta_\ell \rfloor$. This is that $\mathcal{F}$ is *closed under minimization* in every consistent orientation τ of S : that it contains every set $\sigma' \subseteq \tau$ obtained from some $\sigma \subseteq \tau$ in $\mathcal{F}$ by replacing every $\vec{s} \in \sigma$ with some $\vec{s}' \leq \vec{s}$ from τ . Our richness notion in [Definition 3.2](#) will be weaker than this ([\[3, Lemma 4.4\]](#)), but still strong enough to ensure the existence of $\mathcal{F}$ -tangle structure trees.

A separation tree T on an ordered separation system $\vec{S}$ is *ordered* if $|s_v| \leq |s_w|$ whenever v and w are non-leaves of T with $v \leq w$. It is *thoroughly ordered* (in $\vec{S}$) if, for every non-leaf node v , the separation s_v is not oriented by $\lfloor \beta_v \rfloor$ (i.e., $\lfloor \beta_v \rfloor$ contains neither $\vec{s}_v$ nor $\vec{s}_v$) and has minimum order among the separations not oriented by $\lfloor \beta_v \rfloor$. Every thoroughly ordered separation tree is ordered [\[3, Lemma 4.1\]](#).

We say that $\vec{s} \in \vec{S}$ is *weakly eclipsed* by $\vec{r} \in \vec{S}$ if $\vec{r} < \vec{s}$ and $|r| \leq |s|$, and *eclipsed* by $\vec{r}$ if $\vec{r} < \vec{s}$ and $|r| < |s|$. Given any set $\tau \subseteq \vec{S}$, a subset $\sigma \subseteq \tau$ is *efficient* (in τ) if no element of σ is eclipsed by any other element of τ . It is *strongly efficient* if no element of σ is weakly eclipsed by any other element of τ . Our separation tree T is *efficient* if for every leaf ℓ the set β_ℓ is efficient in $\lfloor \beta_\ell \rfloor$. Thoroughly ordered separation trees are efficient [\[3, Lemma 4.9\]](#).

Definition 3.2. A set $\mathcal{F}$ is *rich for $\vec{S}$* if every consistent orientation of S that has a subset in $\mathcal{F}$ also has a strongly efficient⁵ subset in $\mathcal{F}$.

We have the following two theorems for the existence of $\mathcal{F}$ -tangle structure trees.

Theorem 3.3 ([\[3, Theorems 4.6, 4.8\]](#)). *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$. Then there exists a thoroughly ordered $\mathcal{F}$ -tangle structure tree of $\vec{S}$. It is unique if the order function on S is injective.*

Note that the uniqueness part of [Theorem 3.3](#) does not require the existence part: an injectively ordered separation system clearly has at most one thoroughly ordered $\mathcal{F}$ -tangle structure tree, regardless of what $\mathcal{F}$ is. In particular, $\mathcal{F}$ need not be standard or rich. But we know from [\[3, Theorem 4.8\]](#) that $\mathcal{F}$, assuming it is standard, will be rich for $\vec{S}$ if the order function on S is injective and an $\mathcal{F}$ -tangle structure tree exists.

Thoroughly ordered structure trees can be large. But there is a way to reduce them. Given an $\mathcal{F}$ -tangle structure tree (T, r, β) , consider an edge vw of T such that w is a child of v . We call vw *necessary* for a tangle leaf $\ell \geq w$ of T if $\beta(vw)$ is a $\leq$ -minimal element of the tangle $\lfloor \beta_\ell \rfloor$, or equivalently, of β_ℓ . We call vw *necessary* for a forbidden leaf $\ell \geq w$ if every subset of β_ℓ in $\mathcal{F}$ contains $\beta(vw)$. A node v of T is *necessary in T* if for every child w of v there exists a leaf $\ell \geq w$ such that vw is necessary for ℓ . If every node of T is necessary, then T is *irreducible*.

It is easy to make a tangle structure tree irreducible by contracting edges at unnecessary nodes and deleting certain subtrees [\[3, Section 5\]](#). This preserves efficiency and being ordered, but not being thoroughly ordered. Our second existence theorem for $\mathcal{F}$ -tangle structure trees gains efficiency and irreducibility over that from [Theorem 3.3](#) at the expense of losing the thoroughness of its ordering:

Theorem 3.4 ([\[3, Theorem 6.1\]](#)). *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$. Then $\vec{S}$ has an efficient and irreducible ordered $\mathcal{F}$ -tangle structure tree.*

⁵in this orientation of S

See [3, Theorem 6.2] for a long list of properties of these $\mathcal{F}$ -tangle structure trees that make them useful for tangle analysis, both in theory and in applications.

4. TANGLE-TREE DUALITY

If we apply [Theorem 3.4](#) to a separation system $\vec{S}$ that has no $\mathcal{F}$ -tangles, it tells us that there is an $\mathcal{F}$ -tangle structure tree all whose leaves are forbidden. Let us call such structure trees $\mathcal{F}$ -trees. Conversely, if there is an $\mathcal{F}$ -tree of $\vec{S}$ then S has no $\mathcal{F}$ -tangle, by [Theorem 3.1](#).

[Theorem 3.4](#) thus implies the following dichotomy:

Theorem 4.1 ([3, Theorem 6.4]). *Let $\vec{S}$ be an ordered separation system, and let $\mathcal{F}$ be standard and rich for $\vec{S}$. Then exactly one of the following assertions holds:*

- (i) *there exists an $\mathcal{F}$ -tangle of S ;*
- (ii) *there exists an $\mathcal{F}$ -tree of $\vec{S}$.*

In the case of (ii), the $\mathcal{F}$ -tree can be chosen to be ordered, irreducible and efficient.

Note that, unlike in all known theorems of this type, the set $\mathcal{F}$ of forbidden subsets need not consist of stars of separations (see [Section 2](#)).

In the special case that $\mathcal{F}$ does consist of stars of separations, we shall see in [Corollary 4.5](#) that – as long as the $\mathcal{F}$ -tree in (ii) is irreducible, which the theorem says we may require – the separations in $\beta(E(T)) \subseteq \vec{S}$ will be nested. In contexts where separation systems describe genuine separations, of a set or of a graph, say, nested sets of separations are known to cut up the structure they separate in a tree-like way – for example, by a tree-decomposition of the graph. In this way, $\mathcal{F}$ -trees impose a tree structure on such a set or graph, which in turn implies that its separations we are considering have no $\mathcal{F}$ -tangle.

This phenomenon is one of the corner stones of classical tangle theory, known as *tangle-tree duality*. Our aim now is to derive such a dichotomy theorem also in our context, as an application of our more comprehensive tangle structure trees. Our version, [Corollary 4.6](#) below, will not imply the classical tangle-tree duality theorem for abstract separation systems [11], because our richness assumptions made for $\mathcal{F}$ are different. But it will offer the same dichotomy. And we shall see in [2] that it implies all known instances of the classical theorem for concrete structures, such as graphs or matroids.

To carry this out formally, let us introduce the terms in which classical tangle-tree duality is cast. Let $\vec{E}(T)$ denote the set of orientations of the edges of a tree T . For an oriented edge $\vec{e} = (t, t')$ of T we call $t =: i(\vec{e})$ its *initial node* and $t' =: t(\vec{e})$ its *terminal node*, and we denote its inverse (t', t) by $\bar{\vec{e}}$. We think of $t(\vec{e})$ as the node that $\vec{e}$ *points towards*. There is a natural partial ordering $\leq$ on $\vec{E}(T)$ in which $\vec{e} > \vec{f}$ if $f \neq e$ and the unique path in T from e to f starts at $t(\vec{e})$ and ends at $i(\vec{f})$. With the involution $*: \vec{e} \mapsto \bar{\vec{e}}$, this makes $(\vec{E}(T), \leq, *)$ into a separation system.

Since every two edges e, f of T have comparable orientations, the elements of $\vec{E}(T)$ are all nested. And for every node t of T , the set

$$\vec{F}_t := \{ \vec{e} \in \vec{E}(T) : t(\vec{e}) = t \}$$

is a star in this separation system.

Now let $(\vec{S}, \leq, *)$ be any other separation system – typically, but not necessarily, consisting of separations of a set or graph. An *S-tree* (T, α) is a tree T together

with a map $\alpha: \vec{E}(T) \rightarrow \vec{S}$ that satisfies $\alpha(\vec{e}) = \alpha(\vec{e})^*$ for every $\vec{e} \in \vec{E}(T)$. It is an S -tree over a set $\mathcal{F}$ if $\alpha(\vec{F}_t) \in \mathcal{F}$ for every $t \in V(T)$.

If α preserves the partial ordering $\leq$ on $\vec{E}(T)$, i.e. if $\vec{e} \leq \vec{f}$ implies $\alpha(\vec{e}) \leq \alpha(\vec{f})$, and the image of α contains no degenerate separations, then the sets $\alpha(\vec{F}_t) \in \mathcal{F}$ will be stars in $\vec{S}$, because the $\vec{F}_t$ are stars in $\vec{E}(T)$. Conversely:

Lemma 4.2 ([5]). *Let (T, α) be an S -tree over a set stars. If α is injective on $\vec{F}_t$ for every node t of T , then $\vec{e} \leq \vec{f}$ implies $\alpha(\vec{e}) \leq \alpha(\vec{f})$ for all $\vec{e}, \vec{f} \in \vec{E}(T)$. In particular, $\alpha(\vec{E}(T))$ is a nested subset of $\vec{S}$.*

*Conversely, for every nested, regular, finite separation system $(\vec{S}, \leq, *)$ there is an S -tree (T, α) over stars such that $\alpha(\vec{E}(T)) = \vec{S}$ and α is injective on the stars $\vec{F}_t$.*

We remark that if α is not injective on some $\vec{F}_t$, then $\alpha(\vec{E}(T))$ can fail to be nested.

Lemma 4.2 motivates the study of S -trees over sets $\mathcal{F}$ of stars (of separations): they make precise, in a more general way than tree-decompositions do, the ‘tree-like way’ in which nested sets of separations cut up any underlying structure.

The existence of an S -tree (T, α) over a set $\mathcal{F}$, not necessarily consisting of stars, precludes the existence of an $\mathcal{F}$ -tangle of S . Indeed, given any orientation τ of S , the set $\alpha^{-1}(\tau)$ is an orientation of the edges of T . Since trees contain no cycles, this orientation will contain $\vec{F}_t$ for some t , so that $\alpha(\vec{F}_t) \subseteq \tau$. But as (T, α) is an S -tree over $\mathcal{F}$ we have $\alpha(\vec{F}_t) \in \mathcal{F}$, so τ is not an $\mathcal{F}$ -tangle.

The converse implication, that there is an S -tree over $\mathcal{F}$ whenever S has no $\mathcal{F}$ -tangle, holds only when $\mathcal{F}$ consists of stars. But assuming this is not enough. In [11], which offers the most general theorem of this type to date, there is the further assumption that $\vec{S}$ must be $\mathcal{F}$ -separable. This is a richness condition on $\mathcal{F}$, different from ours, which requires $\mathcal{F}$ to also contain stars which its elements ‘induce on the sides’ of certain separations in S . The precise notion of $\mathcal{F}$ -separability is quite technical, but immaterial here, so we refer the reader to [10, 11] for details.

The following dichotomy between $\mathcal{F}$ -tangles and S -trees over $\mathcal{F}$ is the classical tangle-tree duality theorem for abstract separation systems:

Theorem 4.3. [11] *Let $\vec{S}$ be a separation system inside a universe $\vec{U}$ of separations, and let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of stars that is standard for $\vec{S}$. If $\vec{S}$ is $\mathcal{F}$ -separable, then exactly one of the following assertions holds:*

- (i) *there exists an $\mathcal{F}$ -tangle of S ;*
- (ii) *there exists an S -tree over $\mathcal{F}$.*

When $\mathcal{F}$ consists of stars, we can compare **Theorem 4.3** with what our more general **Theorem 4.1** says for such $\mathcal{F}$. Both theorems guarantee that if S admits no $\mathcal{F}$ -tangle then it has an easily detectable certificate in the form of a certain tree-like structure: an S -tree over $\mathcal{F}$ in the case of **Theorem 4.3**, and an $\mathcal{F}$ -tree of $\vec{S}$ in the case of our new **Theorem 4.1**.

The following result shows that our **Theorem 4.1** is not only more general, because its $\mathcal{F}$ need not consist of stars, but also stronger than **Theorem 4.3**, because we can obtain S -trees over $\mathcal{F}$ from $\mathcal{F}$ -trees when $\mathcal{F}$ does consist of stars. In fact, every irreducible $\mathcal{F}$ -tree of $\vec{S}$ then has the property that its edge labels in $\vec{S}$ are precisely the edge labels of some S -tree over $\mathcal{F}$:

Theorem 4.4. *Let $\vec{S}$ be a separation system without trivial elements, $\mathcal{F} \subseteq 2^{\vec{S}}$ a set of stars, and (T, r, β) an irreducible $\mathcal{F}$ -tree of $\vec{S}$. Then there exists an S -tree (T', α) over $\mathcal{F}$ and a map $\gamma: V(T') \cup \vec{E}(T') \rightarrow V(T) \cup E(T)$ such that*

- (i) γ is a bijection from the nodes of T' to the leaves of T ;
- (ii) γ is a bijection from the oriented edges of T' to the edges of T ;
- (iii) for every edge e of T' there is a node v_e of T such that $\{\gamma(\vec{e}), \gamma(\overleftarrow{e})\} = E_{v_e}$, and these v_e are distinct for different e ;
- (iv) for every oriented edge $\vec{e}$ of T' , the map γ sends $t(\vec{e})$ to a leaf $\ell > v_e$ of T such that $\gamma(\vec{e})$ is the first edge of the path $v_e T \ell$;
- (v) $\alpha = \beta \circ \gamma$.

In particular, $\beta(E(T)) = \alpha(\vec{E}(T'))$.

Before we prove [Theorem 4.4](#), let us note a corollary for $\mathcal{F}$ -trees that seems remarkable in its own right, quite independently of tangle-tree duality. If $\mathcal{F}$ consists of stars, the edges of any irreducible $\mathcal{F}$ -tree map to nested separations in $\vec{S}$:

Corollary 4.5. *Let $\vec{S}$ be a separation system without trivial separations, and let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a set of stars. Then for every irreducible $\mathcal{F}$ -tree (T, r, β) of $\vec{S}$ the image $\beta(E(T))$ of its edges is a nested set of separations in $\vec{S}$.*

Proof. Our aim is to apply [Lemma 4.2](#) to the S -tree (T', α) provided by [Theorem 4.4](#), to show that $\alpha(\vec{E}(T')) = \beta(E(T))$ is nested. To apply the lemma, we have to check that α is injective on the stars $\vec{F}_t$ in $\vec{E}(T')$.

Given a node t of T' , we have $t = t(\vec{e})$ for all $\vec{e} \in \vec{F}_t$. For $\ell := \gamma(t)$ we therefore have $\gamma(\vec{F}_t) \subseteq E(rT\ell)$ by [Theorem 4.4 \(iv\)](#). In order to show that α is injective on $\vec{F}_t$, it thus suffices by (v) $\alpha = \beta \circ \gamma$ and the injectivity of γ to show that β is injective on $E(rT\ell)$. This, however, follows from the last clause in the definition of separation trees. $\square$

Proof of Theorem 4.4. We start our construction of T' by taking as $V(T')$ any set we can map to the leaves of T by a bijection γ . This ensures (i).

We shall give T' one edge e for every non-leaf node $v =: v_e$ of T . To choose the ends of e in T' , consider the children w_1 and w_2 of v in T .⁶ For each $i \in \{1, 2\}$, since T is irreducible, the edge vw_i is necessary for some leaf $\ell_i \geq w_i$ of T .⁷ Let e join the nodes $\gamma^{-1}(\ell_1)$ and $\gamma^{-1}(\ell_2)$ of T' , and let $\gamma(\vec{e}) := vw_i$ for the orientation $\vec{e}$ of e for which $\gamma(t(\vec{e})) = \ell_i$. This satisfies (ii)–(iv), if we think of T' as a multigraph for now, i.e., allow multiple edges.

However, since v_e is the unique infimum of ℓ_1 and ℓ_2 in T , the pair $\gamma^{-1}(\{\ell_1, \ell_2\})$ of nodes in T' is joined only by the edge e : the graph T' has no parallel edges and no loops. Finally, now that γ is fixed, let $\alpha := \beta \circ \gamma$ as in (v).

It remains to show that T' is a tree, and that (T', α) is over $\mathcal{F}$: as $\alpha(\vec{e})^* = \alpha(\overleftarrow{e})$ by definition of γ and α , it will be clear that (T', α) is an S -tree once we know that T' is a tree.

Let us begin by showing that T' is acyclic. Consider two edges e, f of T' with $t(\vec{e}) = i(\vec{f})$. By (iv), the edges $\gamma(\vec{e})$ and $\gamma(\overleftarrow{f})$ of T at v_e and v_f , respectively, both lie on the path $rT\ell$ for $\ell = \gamma(t(\vec{e})) = \gamma(i(\vec{f}))$. As $e \neq f$, the nodes v_e and v_f of $rT\ell$ are distinct, by (iii), so $v_e < v_f$ or $v_f < v_e$. As $s_u \neq s_v$ for $u < v$ in any separation tree, $\alpha(\vec{e}) = \beta(\gamma(\vec{e}))$ and $\alpha(\overleftarrow{f}) = \beta(\gamma(\overleftarrow{f}))$, compare (v), are thus orientations of distinct separations $s_{v_e} \neq s_{v_f}$ in S .

⁶These are distinct: s_v cannot be degenerate, because vw_1 is necessary for a forbidden leaf $\ell \geq w_1$ of T , so $\beta(vw_1)$ lies in a star $\sigma \in \mathcal{F}$. Stars do not contain degenerate separations.

⁷Our proof will imply that these ℓ_i are unique, but for the definition it suffices to pick any.

Since we chose $\gamma^{-1}(\ell)$ as $t(\vec{e})$ and as $t(\vec{f})$ when we picked the ends of the edges e and f in T' , the edges $\gamma(\vec{e})$ and $\gamma(\vec{f})$ of T are necessary for ℓ : their β -values $\alpha(\vec{e})$ and $\alpha(\vec{f})$ each lie in some $\sigma \in \mathcal{F}$ that is a subset of β_ℓ , and every subset of β_ℓ in $\mathcal{F}$ contains both. Consider any such $\sigma \in \mathcal{F}$. Since σ is a star, we have $\alpha(\vec{e}) \geq \alpha(\vec{f})^* = \alpha(\vec{f})$. As $s_{v_e} \neq s_{v_f}$, the above inequality is strict. The α -values along any oriented path in T' thus strictly decrease. Hence T' contains no (oriented) cycles.

To complete our proof that T' is a tree, it suffices to show that it has one more node than edges. The first of these numbers equals the number of leaves of T ; the second equals the number of non-leaves of T . This latter number is one less than the former, since T is a binary tree.

To show that (T', α) is an S -tree over $\mathcal{F}$, consider any node t of T' ; we have to show that $\alpha(\vec{F}_t) \in \mathcal{F}$. As earlier, all the edges $\vec{e} \in \vec{F}_t$ are such that $\gamma(\vec{e}) \in rT\ell$ for $\ell := \gamma(t)$ is necessary for ℓ , so $\alpha(\vec{F}_t) = \beta(\gamma(\vec{F}_t)) \subseteq \sigma \subseteq \beta_\ell$ for some $\sigma \in \mathcal{F}$. It remains to show that $\alpha(\vec{F}_t) \supseteq \sigma$.

Suppose there exists an $\vec{s} \in \sigma \setminus \alpha(\vec{F}_t)$. As $\vec{s} \in \sigma \subseteq \beta_\ell = \beta(rT\ell)$, we can find an edge vw in $rT\ell$ with $v < w$ such that $\beta(vw) = \vec{s}$. Let $\vec{f} := \gamma^{-1}(vw)$. Then $\alpha(\vec{f}) = \vec{s}$, and $\gamma(\vec{f}) \in E(rT\ell)$ but $\gamma(\vec{f}) \notin \gamma(\vec{F}_t) \subseteq E(rT\ell)$. In particular, $v_f \neq v_e$ for every $\vec{e} \in \vec{F}_t$, and hence $f \neq e$, so neither $\vec{f}$ or $\vec{f}$ lies in $\vec{F}_t$. But note that v_f and all these v_e lie on the path $rT\ell$, so they are comparable in the tree-order on $V(T)$.

Since T' is a tree, either $\vec{f}$ or $\vec{f}$ points towards t in T' . Suppose first that $\vec{f}$ does. Then $\vec{f} \geq \vec{e}$ in $\vec{E}(T')$ for some $\vec{e} \in \vec{F}_t$. As in our acyclicity proof earlier, this implies that $\alpha(\vec{f}) \geq \alpha(\vec{e})$. As both $\vec{s} = \alpha(\vec{f})$ and $\alpha(\vec{e})$ lie in the star σ , we have $\alpha(\vec{f}) \geq \alpha(\vec{e})^*$ too. As v_f and v_e are distinct nodes on $rT\ell$, which implies that $s = s_{v_f} \neq s_{v_e} = \{\alpha(\vec{e}), \alpha(\vec{e})^*\}$ since T is a separation tree, $\vec{s} = \alpha(\vec{f})$ is thus trivial in $\vec{S}$ witnessed by s_{v_e} , a contradiction to our assumptions about $\vec{S}$.

Suppose now that the edge $\vec{f}$ points towards t in T' . Then $\vec{f} \geq \vec{e} \in \vec{F}_t$ and hence $\alpha(\vec{f})^* \geq \alpha(\vec{e})$, while $s = s_{v_f} \neq s_{v_e} = \{\alpha(\vec{e}), \alpha(\vec{e})^*\}$, as earlier. This makes $\vec{s} = \alpha(\vec{f})$ and $\alpha(\vec{e})$ inconsistent elements of β_ℓ , a contradiction to the fact that T is a consistent separation tree. $\square$

Let us discuss briefly what we can say if $\vec{S}$ does have trivial elements. The proof given above stands until we have to prove that the S -tree T' we have constructed is over $\mathcal{F}$. If $\vec{r} \in \vec{S}$ is trivial in $\vec{S}$ witnessed by $s \in S$, this can indeed fail. Indeed, assume that $\vec{s} = \alpha(\vec{e})$ with $\gamma(t(\vec{e})) = \ell$ and $\gamma(t(\vec{e})) = \ell'$. Let v_f be the predecessor of v_e in T , and assume that $\beta(v_f v_e) = \vec{r}$. Then $v_f v_e$ is the first edge of the path in T from v_f to both ℓ and ℓ' , so either of these would be eligible as $\gamma(t(\vec{f}))$, where $\vec{f}$ is the orientation of f with $\alpha(\vec{f}) = \vec{r}$, as long as $v_f v_e$ is necessary for ℓ and ℓ' , respectively. This will be the case if and only if $\vec{r}$ lies in the stars $\sigma \in \mathcal{F}$ or $\sigma' \in \mathcal{F}$ that contain $\vec{s}$ and $\vec{s}$, respectively; let us assume it lies in both. (Recall that $\beta^{-1}(\vec{s})$ and $\beta^{-1}(\vec{s})$ are necessary for ℓ and ℓ' , respectively, so $\vec{s} \in \sigma$ and $\vec{s} \in \sigma'$ for some and for all stars $\sigma \subseteq \beta_\ell$ and $\sigma' \subseteq \beta_{\ell'}$ in $\mathcal{F}$.) However, when we constructed T' , we chose only one of ℓ and ℓ' as $\gamma(t(\vec{f}))$, say ℓ . We then have $\vec{r} \in \sigma' \setminus \alpha(\vec{F}_{t'})$ for $t' := \gamma^{-1}(\ell')$.

However, we can easily mend this: just add a new leaf t to T' adjacent to t' , with $\alpha(tt') := \vec{r}$. We shall then need that $\{\vec{r}\} \in \mathcal{F}$, but this will be the case if we assume that $\mathcal{F}$ is standard, as we have to in order to obtain T in the first place. More generally, we can add to T' such new leaves at all nodes t' whose $\alpha(\vec{F}_{t'})$ fails to contain a trivial separation $\vec{r}$ whose triviality is witnessed by some s with $\vec{s} \in \alpha(\vec{F}_{t'})$. The

modified tree will no longer satisfy the detailed clauses in [Theorem 4.4](#), but it will be an S -tree over $\mathcal{F}$ all whose non-trivial labels $\alpha(\vec{s})$ are obtained as in our proof.

Note that the ‘tree-likeness’ in S -trees over $\mathcal{F}$ and in $\mathcal{F}$ -trees, respectively, refers to very different structures: while S -trees over $\mathcal{F}$ in [Theorem 4.3](#) exhibit tree-likeness in the (object divided by the) separation system $\vec{S}$ itself – for example, they define tree-decompositions with parts specified by $\mathcal{F}$ if S consists of separations of a graph – the $\mathcal{F}$ -trees in [Theorem 4.1](#) are ‘tree-like’ only in the formal sense of a decision-tree for how to orient the elements of S . It is all the more remarkable, therefore, that the structuring trees found in S by [Theorem 4.3](#) can be recovered from the $\mathcal{F}$ -trees in [Theorem 4.1](#), as shown in [Theorem 4.4](#). We thus have a new tangle-tree duality theorem: one that offers the same dichotomy as [Theorem 4.3](#), but with a different premise.

To state this in optimal form, we need one more definition. Given any set $\mathcal{F} \subseteq 2^{\vec{S}}$ and any order function on S , let $\mathcal{F}_{\text{eff}}$ denote the set of all sets $\sigma \in \mathcal{F}$ that are efficient in their closure $[\sigma]$ in $\vec{S}$ (see [Section 3](#)). Note that this definition makes sense independently of any tangle structure trees.

Corollary 4.6. *Let $\vec{S}$ be an ordered separation system without trivial separations. Let $\mathcal{F} \subseteq 2^{\vec{S}}$ be a rich set of stars. Then exactly one of the following assertions holds:*

- (i) *there exists an $\mathcal{F}$ -tangle of S ;*
- (ii) *there exists an S -tree over $\mathcal{F}$.*

In the case of (ii), the S -tree can be chosen to be over $\mathcal{F}_{\text{eff}}$.

Proof. As indicated after [Lemma 4.2](#), the assertions (i) and (ii) cannot both hold. Let us assume that (i) fails, and show (ii). By [Theorem 4.1](#), there is an $\mathcal{F}$ -tree T of $\vec{S}$, which may be chosen irreducible, efficient and ordered. If T is efficient, it will be an $\mathcal{F}_{\text{eff}}$ -tree, since any $\sigma \subseteq \beta_\ell$ from $\mathcal{F}$ is then efficient in $[\beta_\ell]$ and therefore also efficient in $[\sigma] \subseteq [\beta_\ell]$. Now (ii) and the final statement of the corollary follow from [Theorem 4.4](#). $\square$

While [Theorem 4.3](#) and [Corollary 4.6](#) proclaim the same tangle-tree dichotomy, the requirements they make on $\mathcal{F}$ in their premise are different: $\mathcal{F}$ -separability of $\vec{S}$ and richness of $\mathcal{F}$, respectively. In order to derive [Theorem 4.3](#) formally from [Corollary 4.6](#) (in the case that $\mathcal{F}$ consists of stars) we would need to show that if $\vec{S}$ is $\mathcal{F}$ -separable then $\mathcal{F}$ is rich for $\vec{S}$. The condition of $\mathcal{F}$ -separability, however, was tailor-made for the proof of [Theorem 4.3](#). It therefore seems unlikely that we can formally derive from it that $\mathcal{F}$ is rich. However, in [Section 7](#) we shall derive this from a condition that is only slightly stronger than $\mathcal{F}$ -separability. This stronger condition, in fact, holds and is used in all the known applications of [Theorem 4.3](#), rather than $\mathcal{F}$ -separability itself. This will be explored in [\[2\]](#).

5. SUBMODULAR ORDER FUNCTIONS

Submodular order functions are a key ingredient to much of tangle theory. We shall need them in this paper too, both for our tree-of-tangles theorems in [Section 6](#) and to retrieve the essence of the tangle-tree duality theorem from [\[11\]](#) in [Section 7](#).

In both these applications we shall need tangle structure trees that are not just ordered but thoroughly ordered. These exist reliably only when our order function

is injective. If it is not – which often happens in theoretical applications⁸ – we have to tweak it slightly to make it so. We shall prove in this section that we can do this while preserving its submodularity.

We say that an order function o' on a set S *refines* another order function o on S if

$$o(r) < o(s) \Rightarrow o'(r) < o'(s) \quad (\forall r, s \in S).$$

Thus, o' refines o if and only if $o'(r) \leq o'(s)$ implies $o(r) \leq o(s)$. But we do not have the same implication for ' $<$ ', since o' need not be constant on the subsets of S whose elements have a common value under o .

Let $\vec{S}$ be a separation system with an order function o , and let $\mathcal{F}$ be any set. Recall that the $\mathcal{F}$ -tangles of the subsets

$$S_k = \{s \in S \mid o(s) < k\}$$

of S are the $\mathcal{F}$ -tangles in $\vec{S}$ with respect to o .

Lemma 5.1. *If o' refines o , then the tangles in $\vec{S}$ with respect to o are also tangles in $\vec{S}$ with respect to o' .*

Proof. If o' refines o then, for every $k \in \mathbb{R}$ and with S_k defined with respect to o as above, the range $o'(S_k)$ in $\mathbb{R}$ is an initial segment of $o'(S)$. For every k choose $k' \in \mathbb{R}$ larger than every element of $o'(S_k)$ but smaller than every element of $o'(S) \setminus o'(S_k)$. Then S_k can be rewritten as $\{s \in S \mid o'(s) < k'\}$, which puts the $\mathcal{F}$ -tangles of S_k in $\text{tangle}(\vec{S}, \mathcal{F}, o')$. $\square$

Let us consider a separation system $\vec{U}$ that is a universe: one whose partial ordering $\leq$ makes it into a lattice. Any two $\vec{r}, \vec{s} \in \vec{U}$ thus have a supremum $\vec{r} \vee \vec{s}$ and an infimum $\vec{r} \wedge \vec{s}$ in $\vec{U}$. Let us call a function $u: \vec{U} \rightarrow \mathbb{R}$ *submodular* if

$$u(\vec{r} \vee \vec{s}) + u(\vec{r} \wedge \vec{s}) \leq u(\vec{r}) + u(\vec{s})$$

for all $\vec{r}, \vec{s} \in \vec{U}$, and *structurally submodular* if

$$u(\vec{r} \vee \vec{s}) \leq u(\vec{r}) \quad \text{or} \quad u(\vec{r} \wedge \vec{s}) \leq u(\vec{s}) \quad (\dagger)$$

for all $\vec{r}, \vec{s} \in \vec{U}$. Note that submodular functions on $\vec{U}$ are also structurally submodular, but not conversely. Unlike submodularity, the structural submodularity of a real function on $\vec{U}$ is preserved when we compose it with an order isomorphism on $\mathbb{R}$, but it is not preserved by summing functions.

Let us call an order function $||$ on the set U of unoriented separations *submodular* if the function $\vec{s} \mapsto |s|$ it induces on $\vec{U}$ is submodular, and similarly for 'structurally submodular'. As earlier, submodularity of order functions on U is preserved under summing (but not under composition with order isomorphisms), while structural submodularity of order functions on U is preserved under composition with order isomorphisms (but not under summing).

Note also that, given any $k \in \mathbb{R}$, the separation system $\vec{U}_k = \{\vec{s} \in \vec{U} : |s| < k\}$ is submodular, as defined at the start of [Section 2](#), as soon as $||$ is structurally submodular on $\vec{U}$. Conversely, if $\vec{S} \subseteq \vec{U}$ is submodular, one can define a structurally submodular order function on U so that $\vec{S} = \vec{U}_k$. But there need not exist a submodular order function on U with this property; see [1, Example 5.11].

It is not hard to refine a submodular order function to make it injective. Ideally, we would like to do this while preserving its submodularity. We shall prove a little

⁸Order functions used for real-world tangle applications, such as in clustering, tend to be injective. But the standard order function on graph separations, for example, is not.

less than this, which is still good enough for our use of submodularity: that we can refine it to an order function that is still structurally submodular.

Lemma 5.2. *Let $\vec{U}$ be a separation universe with a submodular function u . Then $u': \vec{s} \mapsto u(\vec{s})$ and $w: s \mapsto u(\vec{s}) + u(\vec{s})$ are also submodular order functions on U .*

Proof. Since u is submodular, we also have

$$u'(\vec{s} \vee \vec{r}) + u'(\vec{s} \wedge \vec{r}) = u(\vec{s} \wedge \vec{r}) + u(\vec{s} \vee \vec{r}) \leq u(\vec{s}) + u(\vec{r}) = u'(\vec{s}) + u'(\vec{r}),$$

so u' is submodular. The sum of two submodular functions is clearly submodular, hence so is also $w = u + u'$. $\square$

Lemma 5.3. *Let $\vec{U}$ be a separation system and $\vec{t} \in \vec{U}$. The function $\mathbf{1}_{\vec{t}}: \vec{U} \rightarrow \{0, 1\}$ given by*

$$\mathbf{1}_{\vec{t}}(\vec{s}) = \begin{cases} 0 & \text{if } \vec{s} \leq \vec{t}, \\ 1 & \text{otherwise.} \end{cases}$$

is submodular.

Proof. Let $\vec{r}, \vec{s} \in \vec{U}$. We have to show that

$$\mathbf{1}_{\vec{t}}(\vec{r} \wedge \vec{s}) + \mathbf{1}_{\vec{t}}(\vec{r} \vee \vec{s}) \leq \mathbf{1}_{\vec{t}}(\vec{r}) + \mathbf{1}_{\vec{t}}(\vec{s}).$$

If the right-hand side of the inequality is 2, the inequality holds trivially. So let us assume that $\mathbf{1}_{\vec{t}}(\vec{r}) = 0$. Then $\vec{r} \wedge \vec{s} \leq \vec{r} \leq \vec{t}$, so $\mathbf{1}_{\vec{t}}(\vec{r} \wedge \vec{s}) = 0$. Hence the left-hand side of the inequality is at most 1. The inequality thus holds unless the right-hand side is zero, i.e., unless also $\mathbf{1}_{\vec{t}}(\vec{s}) = 0$. But then also $\vec{s} \leq \vec{t}$ and thus $\vec{r} \vee \vec{s} \leq \vec{t}$, so $\mathbf{1}_{\vec{t}}(\vec{r} \vee \vec{s}) = 0$, making the left-hand side zero too. $\square$

It may help to think of the sums in our next lemma as n -ary integer expansions.

Lemma 5.4. *Let $\vec{U}$ be a separation universe, $\iota: \vec{U} \rightarrow \{0, \dots, |\vec{U}| - 1\}$ a bijection, and $n \in \mathbb{N}$. Then*

$$\gamma_n^\iota(\vec{s}) := \sum_{\vec{t} \in \vec{U}, \vec{t} \not\leq \vec{s}} n^{\iota(\vec{t})} = \sum_{\vec{t} \in \vec{U}} n^{\iota(\vec{t})} \mathbf{1}_{\vec{t}}(\vec{s}).$$

In particular γ_n^ι is submodular.

Proof. The first part is immediate, as the additional summands in the second sum are all zero. The submodularity of γ_n^ι then follows from [Lemma 5.3](#) and the fact that sums of submodular functions are again submodular. $\square$

Lemma 5.5. *Let $\vec{U}$ be a separation universe with a submodular order function o . Then there exists an injective submodular order function o' on U that refines o .*

Proof. We shall construct a small injective and submodular perturbation $\delta: U \rightarrow \mathbb{R}$, so that $o' := o + \delta$ is injective and refines o .

First choose $\varepsilon > 0$ as follows. Let $\Delta := \{|o(s) - o(r)| : s, r \in U, o(s) \neq o(r)\}$. If $\Delta \neq \emptyset$ set $\varepsilon := \min \Delta > 0$; otherwise set $\varepsilon := 1$. In either case $\varepsilon > 0$.

Fix an arbitrary bijection $\iota: \vec{U} \rightarrow \{0, \dots, |\vec{U}| - 1\}$ and set $\gamma(s) := \gamma_3^\iota(\vec{s}) + \gamma_3^\iota(\vec{s})$. By [Lemma 5.2](#) and [Lemma 5.4](#) the function $\gamma: U \rightarrow \mathbb{R}$ is submodular. We claim γ is injective. Write $\gamma(s)$ in base 3:

$$\gamma(s) = \sum_{k=0}^{|\vec{U}|-1} a_k(s) 3^k, \quad a_k(s) \in \{0, 1, 2\}.$$

By construction we have $a_k(s) \leq 1$ if and only if the separation $\vec{t} = \iota^{-1}(k)$ points towards s . If $\gamma(s) = \gamma(r)$ then their base-3 expansions coincide, hence every $\vec{t} \in \vec{U}$ points towards s if and only if it points towards r . From this we will now deduce $s = r$, showing that γ is injective.

Let $\vec{s}_0$ be any orientation of s . Suppose that, for some $n \in \mathbb{N}$, we have defined an orientation $\vec{s}_n$ of s . Then since $\vec{t} := \vec{s}_n$ points towards s , it also points towards r . So there exists an orientation $\vec{r}_n$ of r with $\vec{s}_n \geq \vec{r}_n$. Analogously, given any orientation $\vec{r}_n$ of r we can find an orientation $\vec{s}_{n+1}$ of s with $\vec{r}_n \geq \vec{s}_{n+1}$. For any orientation $\vec{s}_0$ of s we can thus construct an infinite chain $\vec{s}_0 \geq \vec{r}_0 \geq \vec{s}_1 \geq \dots$ of, alternately, orientations of s and r . As r and s have at most two orientations each, at most two of these inequalities can be strict. Hence for m large enough we have $\vec{s}_m = \vec{r}_m$, and therefore $s = r$ as desired. This completes our proof that γ is injective.

Note that $0 \leq \gamma(s) < 3^{|\vec{U}|}$ for every s . Define

$$\delta(s) := \left(\frac{1}{2}\varepsilon/3^{|\vec{U}|}\right) \gamma(s).$$

Then $0 \leq \delta(s) < \varepsilon/2$ for all s . Note that δ is injective and submodular, because γ is and $\frac{1}{2}\varepsilon/3^{|\vec{U}|} > 0$.

Set $o' := o + \delta$. As a sum of submodular functions, o' is submodular. To see that o' is injective, suppose $o'(r) = o'(s)$. Then

$$|o(r) - o(s)| = |\delta(r) - \delta(s)| \leq |\delta(r)| + |\delta(s)| < \varepsilon.$$

By the choice of ε we must have $o(s) = o(r)$, and hence $\delta(s) = \delta(r)$. Since δ is injective this gives $s = r$. Finally, o' refines o , because whenever $o(r) > o(s)$ we have

$$o'(r) - o'(s) = o(r) - o(s) + \delta(r) - \delta(s) \geq \varepsilon - |\delta(r)| - |\delta(s)| > 0,$$

so $o'(r) > o'(s)$. $\square$

An *enumeration* of a set S is any bijection between S and $\{1, \dots, |S|\}$.

Lemma 5.6. *Let $\vec{U}$ be a separation universe with a submodular order function o . Then there exists a structurally submodular enumeration of U that refines o .*

Proof. Compose the injective submodular refinement o' of o provided by Lemma 5.5 with an order isomorphism φ from its range to $\{1, \dots, |U|\}$. Since o' , being submodular, is also structurally submodular, and structural submodularity is preserved under composition with order isomorphisms, the enumeration $\varphi \circ o'$ of U is structurally submodular, and like o' it refines o . $\square$

6. TREES OF TANGLES

In Section 4 we derived one of the two fundamental types of theorem about tangles from our tangle structure trees: a so-called tangle-tree duality theorem. In this section we shall do the same for the second pillar of tangle theory: its so-called tree-of-tangles theorems.

Let $\vec{S}$ be an ordered separation system. Recall from Section 2 that a separation *distinguishes* two tangles in $\vec{S}$ if they orient it differently. Let us say that it distinguishes them *optimally*⁹ if it distinguishes them but no separation in S of lower

⁹The usual terminology is to say ‘efficiently’ rather than ‘optimally’. We use the latter only in this paper, to avoid confusion with our use of the term ‘efficient’ for subsets of $\vec{S}$.

order does. A nested set $N \subseteq S$ is a *tree*¹⁰ of the tangles of S if every two $\mathcal{F}$ -tangles of S (for some given $\mathcal{F}$) are distinguished optimally by a separation in N . Likewise, N is a *tree of the tangles in $\vec{S}$* if it is nested and every two maximal $\mathcal{F}$ -tangles in $\vec{S}$ are distinguished optimally by a separation in N .

Definition 6.1. Let $\vec{S}$ be a separation system contained in a universe $\vec{U}$. An orientation τ of S is *robust* in $\vec{U}$ if for no $s \in U$ there is a triple $\{\vec{r}, \vec{r} \vee \vec{s}, \vec{r} \vee \vec{s}\}$ in τ with $|\vec{r} \vee \vec{s}|, |\vec{r} \vee \vec{s}| < |r|$.¹¹

These triples will be illustrated in Figure 2. Note that while $\vec{r}$ can lie in such a triple in τ only if it lies in $\vec{S}$, the separation s need not lie in S : it just gives rise to that triple in $\vec{S}$ together with $\vec{r} \in \vec{S}$.

We shall find in every separation system $\vec{S}$ that lives in some universe with an injective submodular order function a tree of all the robust $\mathcal{F}$ -tangles of S , as well as a more comprehensive tree of all the robust $\mathcal{F}$ -tangles in $\vec{S}$. Here, $\mathcal{F}$ has to be standard and rich, as usual, but need not satisfy any further requirements. But to ease our terminology, we shall encode the robustness assumption in $\mathcal{F}$, by assuming that $\mathcal{F}$ contains all the triples from Definition 6.1. We can then simply speak of (trees of) $\mathcal{F}$ -tangles rather than of robust $\mathcal{F}$ -tangles.

The most comprehensive tree-of-tangles theorems known so far are for $\mathcal{F}$ -tangles that are not only robust but also *profiles*: consistent orientations of a separation system $\vec{S}$ that contain no triple of the form $\{\vec{r}, \vec{s}, \vec{r} \vee \vec{s}\}$, where the supremum is once more taken in some fixed universe of separations containing $\vec{S}$ [9].

Tangles of graphs are robust profiles [7].¹² So the tree-of-tangles theorems for robust profiles of abstract separation systems from [9] generalize the tree-of-tangles theorem for graphs of Robertson and Seymour [12], and our results below will further generalize those from [9] to $\mathcal{F}$ -tangles that are robust but need not be profiles.

There is another advance in [9] over [12]: the trees of tangles found in [9] are *canonical*. This means, roughly, that the automorphisms of the given graph or separation system $\vec{S}$ leave the tree of tangles $N \subseteq S$ invariant. This is usually proved by observing that the construction of N depends only on invariants of the graph or separation system rather than, say, of an enumeration of $\vec{S}$ assumed for the sake of the construction, as was the case for the first trees of tangles for graphs in [12]. As a consequence, such a canonical tree of tangles can then be computed by an algorithm that yields the same result regardless of the order in which the elements of the graph or separation system are presented to it.

Our tree-of-tangles theorems for all robust $\mathcal{F}$ -tangles, Theorems 6.2 and 6.10 below, requires that the order function given on S is injective. This assumption makes canonicity, as defined above, trivial: unless $|S| = 1$, the identity will be the only automorphism of $\vec{S}$ as an *ordered* separation system. As we saw in Section 5, we can make a given submodular order function injective without losing its submodularity. But this conversion is not canonical: it depends on choices not determined by the

¹⁰Recall that nested sets of separations of structures such as sets or graphs divide them in a ‘tree-like’ way [5]; hence the ‘tree’ in the name.

¹¹If this definition appears technical, think of orientations of bipartitions of a set as designating the side they point to as ‘big’. Robust orientations of set partitions are then like a finite analogue to ultrafilters: big subsets never split into two small subsets.

¹²Indeed, all tangles of abstract separation systems, as defined in [8], are profiles (see there). If the ambient universe of their separation system is distributive, they are easily seen to be robust. Graph tangles are examples of such abstract tangles in a distributive universe.

invariants of the separation system considered.¹³ If we ignore the order function on our separation system and refer in our notion of canonicity only to the unordered system, as a poset with an order-reversion involution, then the tree of tangles we shall obtain in this section will not be canonical: it will depend on the choice of (injective) order function on S , because the (thoroughly ordered) tangle structure tree from which we construct it depends on it.

The following definition allows us to encode the robustness of our tangles in $\mathcal{F}$. Let $\vec{S}$ be a separation system in a universe $\vec{U}$ of separations. Let

$$\mathcal{R}(\vec{U}) := \{ \{ \vec{r}, \bar{r} \vee \vec{s}, \bar{r} \vee \bar{s} \} \subseteq \vec{U} : r, s \in U \text{ and } |\bar{r} \vee \vec{s}|, |\bar{r} \vee \bar{s}| < |r| \}.$$

These triples are easily identified in [Figure 2](#), for example.

Here is our basic tree-of-tangles theorem for robust $\mathcal{F}$ -tangles with arbitrary $\mathcal{F}$, or equivalently, for all $\mathcal{F}$ -tangles with arbitrary $\mathcal{F}$ that include $\mathcal{R}(\vec{U})$:

Theorem 6.2. *Let $\vec{U}$ be a universe of separations with an injective and structurally submodular order function. Let $k \in \mathbb{R}$ and $\vec{S} := \vec{U}_k$. Let $\mathcal{F}$ be any set that is rich and standard for $\vec{S}$ and includes $\mathcal{R}(\vec{U})$. Let (T, r, β) be the unique thoroughly ordered $\mathcal{F}$ -tangle structure tree for $\vec{S}$ (cf. [Theorem 3.3](#)). Let N be the set of all separations $s_v \in S$ with v a tangle node of T .*

Then N is a tree of the tangles of S ; in particular, it is nested. It consists of all the separations in S that distinguish two $\mathcal{F}$ -tangles of S optimally.

Proof. We shall build up our proof of [Theorem 6.2](#) as a sequence of a few lemmas. Let $\vec{U}$, k , $\vec{S}$, $\mathcal{F}$, (T, r, β) and N be as stated in the theorem, with an injective and structurally submodular order function on U .

The *corners* of two separations $r, s \in S$ are the four separations in U that have orientations of the form $\vec{r} \wedge \vec{s}$. [Figure 2](#) shows a corner of two bipartitions of a set, where $\vec{r} \wedge \vec{s}$ is the intersection of the subsets to which $\vec{r}$ and $\vec{s}$ point.




FIGURE 2. A corner $\vec{t} = \vec{r} \wedge \vec{s}$ of two bipartitions of a set V .

Two corners t and t' of r and s lie on the same side of r if these separations have orientations $\vec{r} \geq \vec{t}, \vec{t}'$. Corners on the same side of one of two crossing separations are *adjacent*; non-adjacent corners are *opposite*. Since our order function on U is structurally submodular, all four ordered pairs (t, t') of opposite corners satisfy $|t| \leq |r|$ or $|t'| \leq |s|$.

In [Figure 2](#), the triples $\{ \vec{r}, \bar{r} \vee \vec{s}, \bar{r} \vee \bar{s} \}$ from the definition of $\mathcal{R}(\vec{U})$ consist of $\vec{r}$ and the two corners on the right-hand side of r both oriented towards r , the last element of the triple being $\vec{t} = \bar{r} \vee \bar{s}$.

¹³In our case, this was the choice of the bijection ι in the proof of [Lemma 5.5](#).

Lemma 6.3. *Any two crossing separations $r, s \in S$ have two adjacent corners that either lie on the same side of r and have order $< |r|$ or lie on the same side of s and have order $< |s|$.*

Proof. Let $|r| \leq |s|$. If all four corners have order $< |s|$ we are done; let t_1 be a corner of order $\geq |s|$. As r and s cross we have $r \neq s \neq t_1$, so in fact $|r| < |s| < |t_1|$.

By structural submodularity, the corner t_2 opposite t_1 has order $\leq |r|$, and hence order $< |r|$, since $t_2 \neq r$ as earlier. If the corner t_3 on the same side of r as t_2 also has order $< |r|$ we are done, so we assume not. Then $|t_3| > |r|$, and its opposite corner t_4 satisfies $|t_4| < |s|$. Since t_4 lies on the same side of s as t_2 , and $|t_2| < |r| < |s|$, this completes the proof. $\square$

Lemma 6.4. (i) *If $s \in S$ distinguishes two tangles τ, τ' of S optimally then $s = s_v$, where v is the tangle node of T that is the infimum in T of the tangle leaves $\ell = \ell(\tau)$ and $\ell' = \ell(\tau')$.¹⁴*

(ii) *Conversely, if the tangle node v of T is the infimum in T of two tangle leaves $\ell = \ell(\tau)$ and $\ell' = \ell(\tau')$, then s_v distinguishes τ and τ' optimally.*

Proof. (i) By definition of v , the tangles $\tau = [\beta_\ell]$ and $\tau' = [\beta_{\ell'}]$ both include $[\beta_v]$, but they differ on s_v . Since s distinguishes τ and τ' optimally, we thus have $|s| \leq |s_v|$, while s is not oriented by $[\beta_v]$. As T is thoroughly ordered, this implies $|s| \geq |s_v|$, and hence $s = s_v$ since $|\cdot|$ is injective.

(ii) The separation s_v distinguishes τ from τ' by definition of v . Since T is thoroughly ordered, any $s \in S$ with $|s| < |s_v|$ has an orientation $\vec{s}$ in $[\beta_v]$, which is a subset of both $\tau = [\beta_\ell]$ and $\tau' = [\beta_{\ell'}]$. Hence $\tau(s) = \tau'(s) = \vec{s}$, so s does not distinguish τ from τ' . $\square$

Lemma 6.5. *N consists of all the separations in S that distinguish two tangles of S optimally.*

Proof. Immediate from Lemma 6.4 and the definition of N . $\square$

Let us call a node v of T *critical* if s_v has an orientation $\vec{s}_v$ that is either co-trivial in $\vec{S}$ or such that $[\beta_v] \cup \{\vec{s}_v\}$ contains a triple from $\mathcal{R}(\vec{U})$.

Lemma 6.6. *Tangle nodes of T are never critical.*

Proof. Suppose T has a critical tangle node v . Let $\vec{s}_v$ be as in the definition of its criticality. Let w be the successor of v with $\beta(vw) = \vec{s}_v$, and let $\ell \geq w$ be a tangle leaf of T .

Then $\vec{s}_v$ lies in the $\mathcal{F}$ -tangle $[\beta_\ell]$. As $\mathcal{F}$ is standard, $\vec{s}_v$ cannot be co-trivial. Hence $[\beta_v] \cup \{\vec{s}_v\}$ contains a triple from $\mathcal{R}(\vec{U}) \subseteq \mathcal{F}$. As $[\beta_v] \cup \{\vec{s}_v\} \subseteq [\beta_\ell]$, this means that ℓ is not a tangle leaf, contrary to its definition. $\square$

Lemma 6.7. *Let u, v be non-leaf nodes of T . If s_u and s_v cross, then either u or v is critical.*

Proof. By Lemma 6.3, we can find $r \in \{s_u, s_v\}$ such that two corners t, t' of s_u and s_v lie on the same side of r and have order $< |r|$. Let $\vec{r}$ orient r away from these corners. Then these have orientations $\vec{t}, \vec{t}' \leq \vec{r}$. Note that $t, t' \in S$, since $|t|, |t'| < |r|$ and $r \in S = U_k$, as well as $\{\vec{r}, \vec{t}, \vec{t}'\} \in \mathcal{R}(\vec{U})$.

¹⁴See Theorem 3.1 for the definition of $\ell(\tau)$ etc.

Let $r = s_v$, say, and let w be the child of v for which $\beta(vw) = \vec{r}$. If $\vec{r}$ is co-trivial in $\vec{S}$, then v is critical as desired, so we assume that $\vec{r}$ is not co-trivial. Since T is an $\mathcal{F}$ -tangle structure tree and v is not a leaf, the set β_v has no subset in $\mathcal{F}$. As $\mathcal{F}$ is standard, this means that β_v , and hence also $\beta_w = \beta_v \cup \{\vec{r}\}$, has no co-trivial elements.

Since t and t' have order $< |r|$ but $s_v = r$, our assumption that T is thoroughly ordered implies that t and t' were ineligible for selection as s_v : they must have orientations in $[\beta_v] \subseteq [\beta_w]$. Now $[\beta_w]$ is consistent by Lemma 2.1, because β_w is consistent by definition of T and has no co-trivial elements. But the only orientations of t and t' that are consistent with $\vec{r} \in \beta_w$ are $\vec{t}$ and $\vec{t}'$, since $\vec{t}, \vec{t}' \leq \vec{r}$. Thus, $\vec{t}, \vec{t}' \in [\beta_v]$.

This makes $\{\vec{r}, \vec{t}, \vec{t}'\} \in \mathcal{R}(\vec{U})$ a witness of the fact that v is critical. $\square$

Lemma 6.8. *N is a tree of the tangles of S . In particular, N is nested.*

Proof. N is nested by Lemmas 6.6 and 6.7. By Lemma 6.5, this makes it a tree of the tangles of S . $\square$

This completes our proof of Theorem 6.2. $\square$

The proof of Theorem 6.2 starts from the premise that we have a thoroughly ordered tangle structure tree for our given separation system. Theorem 3.3 provides this when $\mathcal{F}$ is standard and rich. Both these properties are required in the premise of the theorem; but the richness of $\mathcal{F}$ is never used in the proof. However it follows from the existence of the structure tree that the proof does need, combined with the injectivity of the order function on S [3, Theorem 4.8]. So we may as well require it.

The tangle structure tree in which we find our tree of tangles does not, however, have to be thoroughly ordered: we could work in any tree obtained from the unique thoroughly ordered structure tree provided by Theorem 3.3 by reducing it as in the proof of Theorem 3.4 given in [3]. This is because tangle nodes are never deleted in the reduction process, remain tangle nodes, and their associated separations s_v remain optimal distinguishers of any two tangles such that v is the infimum in T of their tangle leaves. See [3, Section 5] for more analysis.

Theorem 6.2 is our basic tree-of-tangles theorem for robust $\mathcal{F}$ -tangles. It applies only to the tangles of a fixed separation system $\vec{S} = \vec{U}_k$ inside some universe $\vec{U}$. The classical tree-of-tangles theorem for profiles [9], however, is more comprehensive (in its more restricted context): it finds a canonical tree of tangles for all the (maximal) tangles in $\vec{S}$, regardless of their order.¹⁵

We can extend Theorem 6.2 to yield such a more general tree of tangles in $\vec{S}$, too. Our definition of tangle structure trees in this context remains the same as before, except that the notion of tangle leaves is extended to capture all the tangles in $\vec{S}$:

Definition 6.9. An $\mathcal{F}$ -tangle structure tree in $\vec{S}$ is a consistent separation tree T on $\vec{S}$ in which every leaf is either a tangle leaf or forbidden, and for every non-leaf v the set β_v has no subset in $\mathcal{F}$.

Here, a leaf ℓ of T is a *tangle leaf* if $[\beta_\ell]$ is a maximal $\mathcal{F}$ -tangle in $\vec{S}$. It is *forbidden* if β_ℓ has a subset in $\mathcal{F}$.

Definitions that involve the term ‘tangle leaf’, such as that of *necessary* edges or nodes of T , adapt accordingly: they now refer to ‘tangle leaves’ as defined above.

¹⁵Tangles of S_k are often called ‘tangles in $\vec{S}$ of order k ’.

Here, then, is our more comprehensive tree-of-tangles theorem. Its tree of tangles displays more tangles than [Theorem 6.2](#) does, namely, all the maximal tangles in $\vec{S}$ rather than just the tangles of S (which are also maximal tangles in $\vec{S}$). However, its premise is also stronger. So neither theorem implies the other.

Theorem 6.10. *Let $\vec{S}$ be a universe of separations with an injective and structurally submodular order function. Let $\mathcal{F}$ be any set that includes $\mathcal{R}(\vec{S})$ and is rich and standard for $\vec{S}_k$ for every $k \in \mathbb{R}$. Then there exists a tree of the tangles in $\vec{S}$.*

Proof. We start by observing that the thoroughly ordered $\mathcal{F}$ -tangle structure trees (T_k, r_0, β^k) provided by [Theorem 6.2](#) for the separation systems $\vec{S}_k \subseteq \vec{S}$ are unique by [Theorem 3.3](#). And they are naturally nested: growing from the same root r_0 , they satisfy $T_i \subseteq T_j$ and $\beta^i \subseteq \beta^j$ for all $i < j$.

Let n be large enough that $S = S_n$, and let $T := T_n$. Then every non-leaf node v of T is a leaf of some T_k : just choose k big enough that $v \in T_k$, but small enough that $|s_v| \geq k$. As v is a non-leaf node of T , the fact that (T_n, r_0, β^n) is an $\mathcal{F}$ -tangle structure tree implies that $\beta_v^n \supseteq \beta_v^k$ has no subset in $\mathcal{F}$. Thus, v is a tangle leaf of T_k in the usual sense of [Section 3](#).

Let T' be obtained from T by deleting any pairs ℓ, ℓ' of forbidden leaves of T that are children of the same node, and let $\beta := \beta^n \upharpoonright E(T')$. Then every leaf of T' is either a forbidden leaf of T , or a tangle leaf ℓ of T_k for some k . In the latter case, its associated tangle τ of S_k is a maximal tangle in $\vec{S}$: it cannot extend to a tangle $\tau' \supsetneq \tau$ of any $S_{k'}$. Indeed, this τ' would be associated with a tangle leaf $\ell' >_r \ell$ of $T_{k'}$, so τ' would include either $\beta_\ell \cup \{\vec{s}_\ell\}$ or $\beta_\ell \cup \{\vec{s}_\ell\}$. This cannot happen for any $\mathcal{F}$ -tangle τ' , as both these sets have a subset in $\mathcal{F}$ by the choice of ℓ .

Thus, every leaf ℓ of T' is either forbidden (and also a leaf of T), or it is a tangle leaf in our new sense that $[\beta_\ell]$ is a maximal $\mathcal{F}$ -tangle in $\vec{S}$. Conversely, for every maximal tangle τ in $\vec{S}$ there is a leaf ℓ of T' such that $[\beta_\ell] = \tau$. But, unlike in T , for every non-leaf node v of T' there exists a tangle leaf $\ell > v$ of T' .

Let V be the set of all non-leaf nodes of T' that are not the parent of a forbidden leaf of T . This V is precisely the set of all tangle nodes of T' . Let $N = \{s_v \mid v \in V\}$. We shall prove that N is our desired tree of tangles: that it is nested, and that it consists of precisely the separations in S that distinguish two maximal $\mathcal{F}$ -tangles in $\vec{S}$ optimally.

Let us show first that every two maximal tangles in $\vec{S}$, say a tangle τ_i of S_i and a tangle τ_j of S_j with $i \leq j$, are distinguished optimally by some separation in N . Let us apply [Theorem 6.2](#) with $k := i$ to the tangles τ_i and $\tau_j \cap \vec{S}_i \neq \tau_i$ of S_i . By the proof of the theorem, these tangles are distinguished optimally by s_v for v the infimum of their tangle leaves in T_i . This v lies in V , and s_v also distinguishes τ_i and τ_j optimally as tangles in $\vec{S}$, i.e., as partial orientations of S .

Conversely, let us show that every $s_v \in N$ distinguishes some such pair τ_i, τ_j of tangles in $\vec{S}$ optimally. As v lies in V , it has two children in T' , neither of which is a forbidden leaf of T . Since all forbidden leaves of T' are also forbidden leaves of T , the definition of T' implies that v is the infimum of some tangle leaves ℓ_i and ℓ_j of T' associated with tangles τ_i of S_i and τ_j of S_j . By our earlier arguments, these are maximal tangles in $\vec{S}$, and s_v distinguishes them optimally.

To see that N is nested, consider distinct nodes $u, v \in V$ in T . Suppose s_u and s_v cross. By [Lemma 6.3](#), we can find $r \in \{s_u, s_v\}$ such that two corners t, t' of s_u and s_v lie on the same side of r and have order $< |r|$. Let us assume that $r = s_v$. As earlier in our proof, v is a tangle node in some T_i . But now v is critical

in T_i , by the proof of [Lemma 6.7](#). (Note that this proof no longer considers u after identifying v by [Lemma 6.3](#), exactly as we did here. This is important, since in our context u may not be a node of T_i .) This contradicts [Lemma 6.6](#) applied to T_i .¹⁶ $\square$

Let us summarize the details from our proof of [Theorem 6.10](#):

Corollary 6.11. *Let $\vec{S}$ be a universe of separations with an injective and structurally submodular order function. Let $\mathcal{F}$ be any set that includes $\mathcal{R}(\vec{S})$ and is rich and standard for $\vec{S}_k$ for every $k \in \mathbb{R}$. Let T' be obtained from the unique thoroughly ordered $\mathcal{F}$ -tangle structure tree (T, r, β) for $\vec{S}$ by deleting any pairs ℓ, ℓ' of forbidden leaves of T that are children of the same node. Let V be the set of tangle nodes of T' . Then $N = \{s_v \mid v \in V\}$ is a tree of the tangles in $\vec{S}$. It is nested and consists of all the separations in S that distinguish two maximal $\mathcal{F}$ -tangles in $\vec{S}$ optimally.*

7. DISCOVERING RICHES BY ZOOMING IN: HOW TO MAKE CONSISTENT SUBSETS MORE EFFICIENT

When we compared our new tangle-tree duality results from [Section 4](#) with the classical such theorems, such as [Theorem 4.3](#) [11], we observed two things. One was that our dichotomy [Theorem 4.1](#) is more general than that from [11] since, unlike there, our sets $\mathcal{F}$ of forbidden subsets need not be stars of separations. Due to this greater generality, however, our witnesses for the non-existence of tangles, though ‘tree-like’ in the sense that they are $\mathcal{F}$ -trees, do not impose a tree structure on whatever our separation system separates (e.g., a graph or dataset), as [Theorem 4.3](#) from [11] does.

For the case that $\mathcal{F}$ does consist of stars we went on to prove that our results imply a tangle-tree duality theorem just like [Theorem 4.3](#), albeit from a different premise. At the end of [Section 4](#) we remarked that, since our premise (that $\mathcal{F}$ is rich) differs so fundamentally from the premise in [11] (that $\vec{S}$ is $\mathcal{F}$ -separable), it is unlikely that our results could imply those of [11] formally.

Our aim in this section is to show that our premise (and hence, tangle-tree duality) does, however, follow from a slight strengthening of $\mathcal{F}$ -separability, one that is actually verified and used when the duality theorem of [11] is applied [10]. As a result, we shall be able to deduce all the applications of [11] to concrete structures of interest [10] from our results. This will be done in [2].

Assuming a structurally submodular injective order function on $\vec{S}$, we shall prove that $\mathcal{F}$ is rich as soon as it is *closed under shifting*. We can then apply our [Corollary 4.6](#) to obtain the same tangle-tree dichotomy as offered by [Theorem 4.3](#), which also makes this assumption about $\mathcal{F}$. In fact, our notion of ‘closed under shifting’ will be weaker than that in [11], making our result stronger. While [11] makes no formal submodularity assumption, it assumes a consequence of it (separability of $\vec{S}$), which in practice is always established by assuming submodularity [10, Lemma 3.4].

Our strategy for proving, from assumptions similar to those made in [11], that some given $\mathcal{F}$ is rich, is as follows. Recall that $\mathcal{F}$ is *rich* if every consistent orientation τ of S that has a subset in $\mathcal{F}$ also has an efficient such subset: one that is ‘maximally focused’ in the sense that none of its elements is eclipsed by another separation in τ . Our strategy will be to start from an arbitrary set $\sigma \in \mathcal{F}$, and

¹⁶We cannot apply the two lemmas to T' directly because T' , unlike T and T_i , need not be a tangle structure tree of any $\vec{S}'$ with $S' \subseteq S$.

then to make it more focused by iteratively ‘zooming in’ until it is efficient.

For intuition, think of any consistent orientation τ of $\vec{S}$ as pointing towards some place in a dataset V that S may be separating. If τ is a tangle, this could be a cluster; if not, it could be the (possibly empty) interior $\bigcap \sigma$ of a star $\sigma \in \mathcal{F}$. For two oriented separations $\vec{r}, \vec{s}$ in τ , think of $\vec{s} > \vec{r}$ as expressing that $\vec{s}$ points towards r and, beyond it, to whatever place $\vec{r}$ points to in V . Thus, while $\vec{r}$ and $\vec{s}$ point to the same place, $\vec{r}$ is closer to it. Hence if $\vec{r}$ eclipses $\vec{s}$, then replacing $\vec{s}$ with $\vec{r}$ in some subset σ of τ – e.g., one in $\mathcal{F}$, or a set of elements whose closure is a tangle of S – makes σ more focused on the place to which τ points.

Our strategy for making a given set $\sigma \in \mathcal{F}$ more efficient will thus be to replace, iteratively, an element of σ with another element of τ that eclipses it. The rest of σ will either be left unchanged (at the start of this section) or ‘shifted’ to one side of the eclipsing separation, as in [11] (later in this section).

Let us call a collection $\mathcal{F}$ of subsets of $\vec{S}$ *closed under eclipsing* in $\tau \subseteq \vec{S}$ if any set σ' obtained from a subset $\sigma \in \mathcal{F}$ of τ by replacing some $\vec{s} \in \sigma$ with a separation $\vec{r} \in \tau$ that weakly eclipses $\vec{s}$ will also lie in $\mathcal{F}$. Note that σ' will be strictly smaller than σ in the natural extension of our partial ordering $\leq$ on $\vec{S}$ to $2^{\vec{S}}$, in which $\sigma \geq \sigma'$ means that there exists a map $f: \sigma \rightarrow \sigma'$ such that $\vec{s} \geq f(\vec{s})$ for all $\vec{s} \in \sigma$.

Lemma 7.1. *If $\vec{S}$ is an ordered separation system and $\mathcal{F} \subseteq 2^{\vec{S}}$ is closed under eclipsing in every consistent orientation of S , then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. To prove that $\mathcal{F}$ is rich, let τ be any consistent orientation of S that has a subset in $\mathcal{F}$. We claim that every set σ that is minimal in $\mathcal{F} \cap 2^\tau \neq \emptyset$ with respect to our partial ordering on $2^{\vec{S}}$ is strongly efficient in τ .

Indeed, suppose σ is not strongly efficient in τ . Then there exists an $\vec{r} \in \tau$ that weakly eclipses some $\vec{s} \in \sigma$. Let $\sigma' := (\sigma \setminus \{\vec{s}\}) \cup \{\vec{r}\}$. Then $\sigma > \sigma' \in \mathcal{F}$, since $\mathcal{F}$ is closed under eclipsing in τ . This contradicts the minimality of σ , since $\sigma' \subseteq \sigma$. $\square$

We now introduce our second way of making a set $\sigma \in \mathcal{F}$ more focused, which is essentially the shifting operation from [10, 11]. Let $\vec{s} \in \vec{S}$ be non-trivial in $\vec{S}$ and non-degenerate. Write $S_{\leq \vec{s}}$ for the set of all $t \in S$ to which $\vec{s}$ points, those with an orientation $\vec{t} \leq \vec{s}$. Given any $\vec{r} \leq \vec{s}$ in $\vec{S}$, define the $\vec{S}_{\leq \vec{s}} \rightarrow \vec{U}$ map¹⁷

$$f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) := \begin{cases} \vec{t} \wedge \vec{r} & \text{if } \vec{s} \neq \vec{t} \leq \vec{s}; \\ (\vec{t} \wedge \vec{r})^* & \text{if } \vec{s} \neq \vec{t} \leq \vec{s}. \end{cases}$$

It is easy to see that such *shifting maps* preserve the partial ordering between their arguments; in particular, they map stars to stars [11]. We say that a separation $\vec{r} \leq \vec{s}$ in $\vec{S}$ *emulates* $\vec{s}$ in $\vec{S}$ if our shifting map has its image inside $\vec{S}$, i.e., if $\vec{t} \wedge \vec{r} \in \vec{S}$ for every $\vec{t} \in \vec{S}$ with $\vec{s} \neq \vec{t} \leq \vec{s}$.

Let us show that shifting stars can reduce their *order*, the sum of the orders of their elements under a given order function on our separation system.

Lemma 7.2. *Let $\vec{U}$ be a universe of separations with a structurally submodular order function, and let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Let τ be a consistent orientation of S . Assume that some element $\vec{s}$ of a star $\sigma \subseteq \tau$ is non-trivial in $\vec{S}$ and eclipsed by some $\vec{r} \in \tau$. Then $\vec{r}$ can be chosen so that it emulates $\vec{s}$ in $\vec{S}$ and $\sigma' := f \downarrow_{\vec{r}}^{\vec{s}}(\sigma)$ is a star in τ of lower order than σ .*

¹⁷Note that, as $\vec{s}$ is non-trivial, $\vec{t} \leq \vec{s}$ and $\vec{t} \leq \vec{s}$ cannot both hold unless $s = t$, in which case $\vec{s} \wedge \vec{r} = \vec{r}$ (rather than its inverse) is chosen as the image of $\vec{s}$, and $\vec{r}$ as the image of $\vec{s}$.

Proof. Choose $\vec{r}$ with $|r|$ minimum, and subject to this maximal under the partial ordering on $\vec{S}$. Let us show first that $\vec{r}$ emulates $\vec{s}$ in $\vec{S}$.

We have to show that $\vec{r} \wedge \vec{t} \in \vec{S}$ for any $\vec{t} \in \vec{S}$ with $\vec{s} \neq \vec{t} \leq \vec{s}$, i.e., that $|\vec{r} \wedge \vec{t}| < \ell$. By the structural submodularity of our order function we have $|\vec{r} \wedge \vec{t}| \leq |t| < \ell$ as desired if $|\vec{r} \vee \vec{t}| > |r|$, so let us assume that $|\vec{r} \vee \vec{t}| \leq |r|$. Then $\vec{r} \vee \vec{t} \in \vec{S}$; let us show that, in fact, $\vec{r} \vee \vec{t} \in \tau$.

As $\vec{r} \vee \vec{t} \geq \vec{r} \in \tau$, the consistency of τ implies $\vec{r} \vee \vec{t} \in \tau$ as desired, unless the unoriented separation underlying $\vec{r} \vee \vec{t}$ is r . If $\vec{r} \vee \vec{t} = \vec{r}$, we are done since $\vec{r} \in \tau$, so let us assume that $\vec{r} = \vec{r} \vee \vec{t} \geq \vec{r}$. Then $\vec{s} \geq \vec{r} \vee \vec{t} = \vec{r}$, since $\vec{s} \geq \vec{r}$ as well as $\vec{s} \geq \vec{t}$ by assumption, which makes $\vec{s}$ trivial witnessed by r , contrary to assumption, unless $s = r$. But if $\vec{s} = \vec{r}$ we have $\vec{r} \vee \vec{t} = \vec{s} \vee \vec{t} = \vec{s} \in \tau$, while if $\vec{s} = \vec{r}$ we have $\vec{r} \vee \vec{t} = \vec{r} = \vec{s} \in \tau$. This completes our proof that $\vec{r} \vee \vec{t} \in \tau$.

As noted, we have $\vec{s} \geq \vec{r} \vee \vec{t}$, as well as $|\vec{r} \vee \vec{t}| \leq |r| < |s|$, since $\vec{r}$ eclipses $\vec{s}$. So $\vec{r} \vee \vec{t}$ eclipses $\vec{s}$, and was thus a candidate for $\vec{r}$. As $|\vec{r} \vee \vec{t}| \leq |r|$, this contradicts the choice of $\vec{r}$ unless $\vec{r} \vee \vec{t} = \vec{r}$. But in that case $\vec{r} \geq \vec{t}$, giving $\vec{r} \wedge \vec{t} = \vec{t} \in \vec{S}$ as desired. This completes our proof that $\vec{r}$ emulates $\vec{s}$.

As shifting preserves the partial ordering on $\vec{S}$, we know that σ' is a star. Let us show that it has lower order than σ . Since $f \downarrow_{\vec{r}}^{\vec{s}}(\vec{s}) = \vec{r}$ and $|r| < |s|$, it suffices to show that $|f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t})| = |\vec{r} \wedge \vec{t}| \leq |t|$ for any other elements $\vec{t}$ of σ . But this follows from structural submodularity and the choice of $\vec{r}$, as earlier, unless $\vec{r} \wedge \vec{t} = \vec{t}$. But in that case we also have $|\vec{r} \wedge \vec{t}| = |t|$ as desired.

It remains to show that $\sigma' \subseteq \tau$. We have $f \downarrow_{\vec{r}}^{\vec{s}}(\vec{s}) = \vec{r} \in \tau$ by explicit assumption about $\vec{r}$. For any $\vec{t} \in \sigma$ other than $\vec{s}$, we have $\vec{s} \geq \vec{t}$, since σ is a star, and hence $f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) = (\vec{t} \wedge \vec{r})^* = \vec{t} \vee \vec{r} \geq \vec{t} \in \tau$ by the definition of $f \downarrow_{\vec{r}}^{\vec{s}}$. This implies $f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) \in \tau$ by the consistency of τ unless $f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) = \vec{t}$, so let us show that this cannot be the case.

If it is, then we have shown that $\vec{t} = f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) \geq \vec{t}$. Then $\vec{s} \geq \vec{t} \geq \vec{t}$, which contradicts our assumption that $\vec{s}$ is non-trivial unless $s = t$. But in that case we have $\vec{s} = \vec{t}$, since $\vec{s} \neq \vec{t}$ as these are distinct elements of σ . But if $\vec{s} = \vec{t}$ then $\vec{r} = f \downarrow_{\vec{r}}^{\vec{s}}(\vec{s}) = f \downarrow_{\vec{r}}^{\vec{s}}(\vec{t}) = \vec{t} = \vec{s}$, which contradicts the fact that $|r| < |s|$. $\square$

Call a set $\mathcal{F} \subseteq 2^{\vec{S}}$ of stars *closed under shifting*¹⁸ in $\tau \subseteq \vec{S}$ if for every $\sigma \in \mathcal{F} \cap 2^{\tau}$ and every $\vec{r} \in \tau$ that weakly eclipses and emulates some non-trivial $\vec{s} \in \sigma$ in $\vec{S}$ we have $f \downarrow_{\vec{r}}^{\vec{s}}(\sigma) \in \mathcal{F}$.

Lemma 7.3. *Let $\vec{U}$ be a universe of separations with a structurally submodular injective order function, and let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Let $\mathcal{F}$ be a set of stars of non-trivial separations in $\vec{S}$ that is closed under shifting in $\vec{S}$.¹⁸ Then $\mathcal{F}$ is rich for $\vec{S}$.*

Proof. To prove that $\mathcal{F}$ is rich, consider any consistent orientation τ of S that includes some $\sigma \in \mathcal{F}$. Choose such a star σ in τ of minimum order. We shall show that σ is strongly efficient in τ .

If not, there exists an $\vec{r} \in \tau$ that weakly eclipses some $\vec{s} \in \sigma$. Since our order function is injective, any such $\vec{r}$ in fact eclipses $\vec{s}$. By Lemma 7.2, we can choose $\vec{r}$ so that it emulates $\vec{s}$ in $\vec{S}$ and $\sigma' := f \downarrow_{\vec{r}}^{\vec{s}}(\sigma)$ is a star in τ of lower order than σ . As $\mathcal{F}$ is closed under shifting we have $\sigma' \in \mathcal{F}$, which contradicts our choice of σ . $\square$

¹⁸Our definition of ‘closed under shifting’ is weaker than that in [10, 11], making our subsequent results stronger. We can even strengthen them further: as is easily checked, we can weaken in both Lemma 7.3 and Theorem 7.4 the assumption that $\mathcal{F}$ is closed under shifting in $\vec{S}$ to assuming that $\mathcal{F}$ is closed under shifting in every consistent orientation of S . We omitted this for simplicity.

Theorem 7.4. *Let $\vec{U}$ be a universe of separations with either a submodular or an injective structurally submodular order function. Let $S = U_\ell$ for some $\ell \in \mathbb{R}$. Assume that $\vec{S}$ has no trivial elements. Let $\mathcal{F}$ be a set of stars of separations in $\vec{S}$ that is closed under shifting in $\vec{S}$.¹⁸ Then exactly one of the following assertions holds:*

- (i) *there exists an $\mathcal{F}$ -tangle of S ;*
- (ii) *there exists an S -tree over $\mathcal{F}$.*

Proof. Recall from the proof of Lemma 5.1 that if S has the form of $S = U_\ell$ with respect to a given order function o , it also has this form with respect to any refinement of o . By Lemma 5.6 we may therefore assume that our order function on U is injective and structurally submodular. By Lemma 7.3, $\mathcal{F}$ is rich for $\vec{S}$. The theorem now follows from Corollary 4.6. $\square$

While Theorem 7.4 is probably to be the most widely applicable tangle-tree duality theorem to date, it may appear restrictive that S must have the form of U_ℓ , and that we need a submodular order function on U . The original tangle-tree duality theorem from [11], Theorem 4.3, makes no such assumptions. However, it makes another assumption, one not made in Theorem 7.4: that $\vec{S}$ is $\mathcal{F}$ -separable.

This is not a coincidence: $\mathcal{F}$ -separability is essentially a technical summary of exactly those consequences of our two more natural assumptions, submodularity and $S = U_\ell$, that are technically needed in the proof of Theorem 4.3.

Finally, the premise of Theorem 7.4 differs from that of Theorem 4.3 in that it requires $\vec{S}$ not to have trivial elements. We need this for our conversion of $\mathcal{F}$ -trees into S -trees over $\mathcal{F}$ in Theorem 4.4. But it entails no loss of generality: the trivial separations of any separation system lie in all consistent orientations, so we can just delete them from our given system and add them back later if desired [4, 5].

REFERENCES

- [1] Sandra Albrechtsen, *Refining trees of tangles in abstract separation systems: Inessential parts (extended version)*, Combinatorial Theory **4**(1) (2024; arXiv:2302.01808v1).
- [2] Hanno von Bergen, *Tangle structure tree applications*, in preparation.
- [3] Hanno von Bergen and Reinhard Diestel, *Tangle structure trees*, arXiv:2601.15870.
- [4] R. Diestel, *Abstract separation systems*, Order **35** (2018), 157–170; arXiv:1406.3797.
- [5] ———, *Tree sets*, Order **35** (2018), 171–192.
- [6] ———, *Tangles: A structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press, 2024. e-book: tangles-book.com/book/.
- [7] ———, *Graph theory* (6th edition), Springer-Verlag, 2025.
Electronic edition available at <http://diestel-graph-theory.com/>.
- [8] R. Diestel, J. Erde, and D. Weißauer, *Structural submodularity and tangles in abstract separation systems*, J. Comb. Theory Ser. A **167C** (2019), 155–180; arXiv:1805.01439.
- [9] R. Diestel, F. Hundertmark, and S. Lemanczyk, *Profiles of separations: in graphs, matroids, and beyond*, Combinatorica **39** (2019), no. 1, 37–75; arXiv:1110.6207.
- [10] R. Diestel and S. Oum, *Tangle-tree duality in graphs, matroids and beyond*, Combinatorica **39** (2019), 879–910. arXiv:1701.02651.
- [11] ———, *Tangle-tree duality in abstract separation systems*, Advances in Mathematics **377** (2021), 107470; arXiv:1701.02509.
- [12] N. Robertson and P.D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Comb. Theory Ser. B **52** (1991), 153–190.